

VNX INTEGRATIONS FOR OPTIMAL PERFORMANCE

Omar Aboulfotoh
EMC Corporation

Wessam Seleem
EMC Corporation

Maha Atef
EMC Corporation

Mohammed Hashem
EMC Corporation

EMC²

Table of Contents

Executive Summary 3

Preface 4

Differences between 1st generation and 2nd generation VNX 6

First Generation VNX 6

The Next-Generation of VNX 8

VNX integration with VMware 10

VNX integration with Recover Point 15

VNX integration with ATMOS 19

VNX Best Practices for Optimal Performance 22

VNX1 22

VNX2 27

Author Biographies 29

Appendix 30

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect EMC Corporation's views, processes or methodologies.

Executive Summary

Digital information is increasing rapidly at unbelievable rates. It needs to be stored on high tier storage to be available, secured, and accessible 24x7. There are many types of the storage arrays, each having capabilities and features to differentiate it and meet customer requirements.

EMC offers storage technologies that maintain data availability, along with a number of other features. One EMC product, VNX®, is one of the classic, sophisticated, and dominant storage solutions in the current data center environment.

This article will explain what VNX and its features and will look at the differences between VNX1 and VNX2, the new generation of this product that contains many new features to facilititate many tasks for the customer and enhance performance.

EMC VNX is very flexible product. It can be integrated with different EMC or non-EMC products. This integration helps to provide many new features, facilitate customer tasks, and maintain data availability and accuracy.

This article will clarify many important points in the integration of EMC VNX and other products like VMware, RecoverPoint, and ATMOS®. We will illustrate how VNX can integrate with all of these products, the benefits of this integration, and how to avoid being stuck on limitations by employing best practices to optimize performance so that all customers realize high performance from their products. It will also help the Sales Team provide customers with a package of well-integrated, complete solutions of EMC storage products operating in rhythm as one product.

Preface

Whether in a shrinking or expanding economy or mature or emerging markets, industries search for profit and IT is a major component of aiding that search. However, key IT challenges still exist and classical approaches to build and manage IT infrastructure no longer make economic sense. There are four basic challenges for any IT decision maker:

1. **Overcome flat budgets**
   Budgets remain flat to nearly up and are not growing nearly fast enough to meet IT demands using conventional approaches.
2. **Manage escalating complexity**
   Companies struggle to manage increasing complexity and are searching for new ways to keep it easy to manage.
3. **Cope with rapidly expanding data growth**
   Data is increasing at an unbelievable rate, requiring solutions to keep up to date.
4. **Meet increased business demands and requirements**
   Business models can be changed many times. Also, competitive pressure and other factors are putting increased demand on IT operations.

EMC designs solutions to overcome these challenges and make the life easier for customers and users. The EMC VNX Family meets these challenges with fundamental design approaches backed by leading innovation and creative technologies.

EMC VNX hardware and software solutions are very simple to provision, efficient, affordable for any budget, and powerful enough to handle the demands.

The Midrange VNX series – Storage Area Network (SAN) product, CLARiiON® and a Network Attached Storage (NAS) product, Celerra® –is combination of unified SAN, NAS, scalability, availability, and performance.



EMC Unisphere® software is a very simple tool to manage the VNX, provideing a common unified management capability for the VNX family, CLARiiON, and Celerra. Also, simplifies and automates other common storage tasks such as replication and backup operations

VNX can be integrated with many products, i.e. RecoverPoint, Symmertirx®, and VMware which facilitate many tasks and enhance performance.

Differences between 1st generation and 2nd generation VNX

First Generation VNX

VNX is Unified Storage that combines many storage protocols that are needed in today's IT industry. VNX supports Fibre Channel, iSCSI, CIFS, and NFS protocols. The first generation of VNX series includes VNX5100 (FC block only), VNX5300, VNX5500, VNX5700, and VNX7500 serving both Block and File components in Unified Platform.

The table below shows the main differences between each model.

		VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
FAST Cache Max Size		100	500	1000	1500	2100
Max. Disks		75	125	250	500	1000
IO Slots		0	4 – 8	4 - 13	6 – 18	6 - 42
File	X-Blades	NA	1 – 2	1, 2 or 3	2, 3 or 4	2 – 8
	System Memory	NA	6 GB	12 GB	12 GB	24 GB
	Protocols	NA	CIFS, NFS and pNFS			
Block	SPs	2				
	System Memory	4 GB	8 GB	12 GB	18 GB	24 GB
	Protocols	FC	FC, iSCSI and FCoE			

VNX provides many features and benefits that improve overall performance. The table below highlights some of the features that improve performance and efficiency.

Optimized for Virtualization	VNX is optimized for virtualization and is an ideal choice for VMware. EMC provides several APIs delivered on VMware to allow VNX to be fully automated and optimized for virtualization.
FAST Cache	Optional Cache of Flash drives providing low latency with high IO.
FAST VP	Automatically locate the right data on the right place, ie place the frequently accessed data on the most high performance disks.
Compression and Deduplication	Reducing the cost by reducing the disk space used.
Object Storage support via Atmos VE	Atoms VE with VNX provides an easy and powerful way to build Object data store.

The Next-Generation of VNX

The new generation of VNX introduces new features and enhancements over the first generation, including Multicore Cache, FAST® VP enhancements, larger FAST Cache, block deduplication, and support for SMB3.

The Next-generation VNX series includes VNX5200, VNX5400, VNX5800, VNX7600, and VNX8000.

The table below shows the main differences between each model.

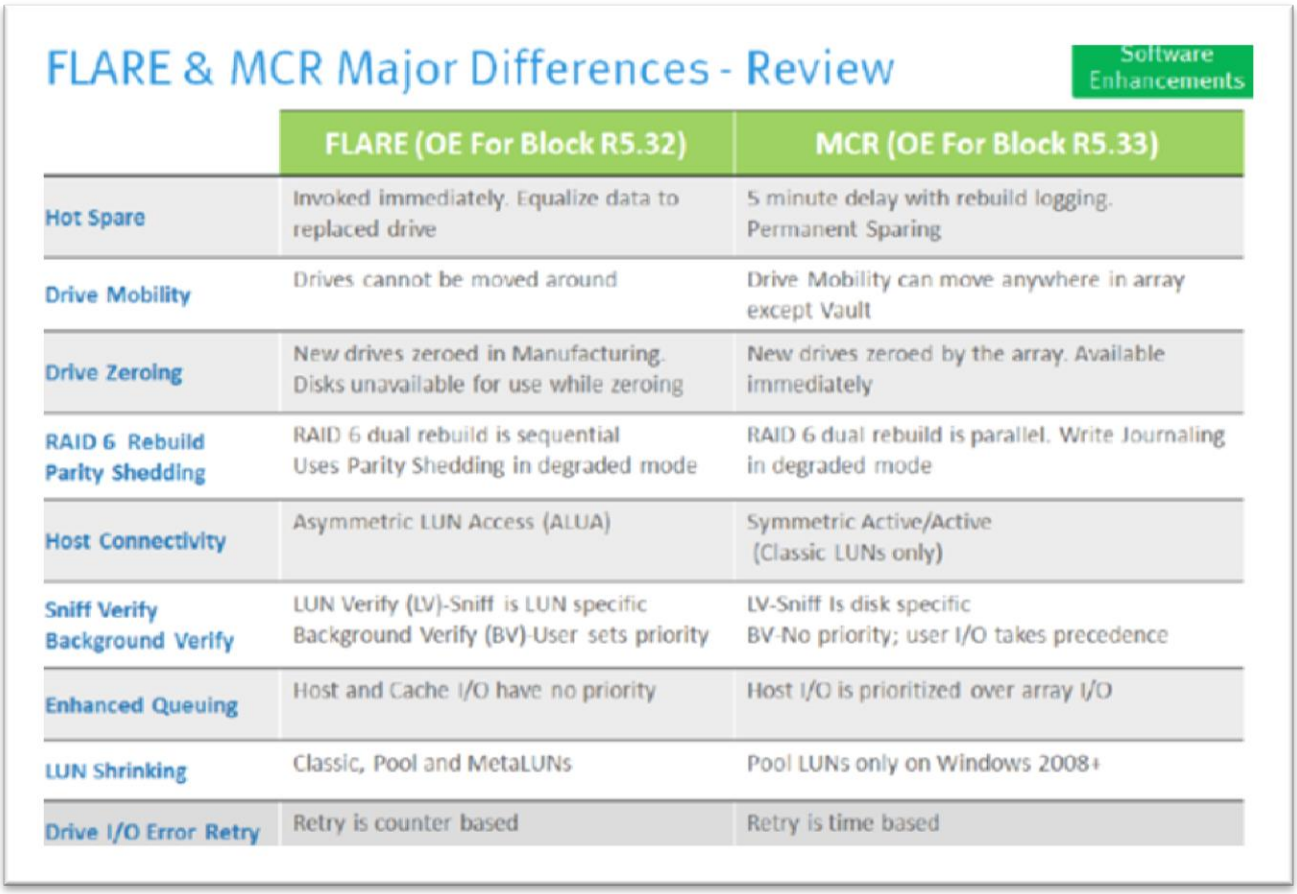
		VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
FAST Cache Max Size		600	1000	2000	3000	4200	4200
Max. Disks		125	250	500	750	1000	1500
File	X-Blades	1 or 2	1 or 2	1, 2 or 3	2, 3 or 4	2 – 4	2 – 8
	System Memory	6 GB	6 GB	12 GB	12 GB	24 GB	24 GB
	IO per Blade	3	3	3	4	4	6
	Protocols	CIFS, NFS and pNFS					
Block	SPs	2					
	System Memory	16 GB	16 GB	24 GB	32 GB	64 GB	128 GB
	Protocols	FC		FC, iSCSI and FCoE			

New and Enhanced Features

New features and enhancements in the next generation of VNX include:

- **Multicore Optimization (MCx):** New technology designed to leverage all of Intel's cores and sockets so that processes take full advantage of multicore CPUs, significantly improving performance.
- **Multicore Cache:** Automatically assign the write cache value for the read according to system needs.
- **Enhanced FAST VP:** MCx now relocate data between tiers using 256 MB slices, increasing disk space efficiency and decreasing relocation impact on the system.
- **Block Level Deduplication:** A new feature in MCx, deduplication can now be enabled on a set of LUNs reducing the space used on LUNs by removing similar data chucks.
- **Symmetric LUN Access:** MCx now uses symmetrical Active-Active connectivity to backend LUNs. Traditional LUNs now can be seen from both storage processors.
- **One-button Shutdown:** New feature provides the ability to safely power down the array with only one click using Unisphere.
- **SMB3 support:** Next-Generation VNX supports SMB3 protocol as part of Windows 8 and Windows 2012.

Major differences between VNX1 and VNX2



FLARE & MCR Major Differences - Review

Software Enhancements

	FLARE (OE For Block R5.32)	MCR (OE For Block R5.33)
Hot Spare	Invoked immediately. Equalize data to replaced drive	5 minute delay with rebuild logging. Permanent Sparing
Drive Mobility	Drives cannot be moved around	Drive Mobility can move anywhere in array except Vault
Drive Zeroing	New drives zeroed in Manufacturing. Disks unavailable for use while zeroing	New drives zeroed by the array. Available immediately
RAID 6 Rebuild Parity Shedding	RAID 6 dual rebuild is sequential Uses Parity Shedding in degraded mode	RAID 6 dual rebuild is parallel. Write Journaling in degraded mode
Host Connectivity	Asymmetric LUN Access (ALUA)	Symmetric Active/Active (Classic LUNs only)
Sniff Verify Background Verify	LUN Verify (LV)-Sniff is LUN specific Background Verify (BV)-User sets priority	LV-Sniff Is disk specific BV-No priority; user I/O takes precedence
Enhanced Queuing	Host and Cache I/O have no priority	Host I/O is prioritized over array I/O
LUN Shrinking	Classic, Pool and MetaLUNs	Pool LUNs only on Windows 2008+
Drive I/O Error Retry	Retry is counter based	Retry is time based

VNX integration with VMware

VMware takes advantage of the array-storage capabilities provided by VNX unified systems via block and file level access. Use of the iSCSI, Fibre Channel, FCoE, and NFS storage protocols provide standard TCP/IP and Fibre Channel network services, as well as storage array features and capabilities which deliver a complete multi-protocol foundation for a VMware vSphere virtual data center.

EMC Virtual Storage Integrator (VSI) for VMware vCenter with VNX

- To view the plugins installed, path management and Unified Storage management can be installed into vSphere.

Access to VNX:

- Storage protocol (Block and File).
- Figure management console information.
- VNX can provide dedicated pools to be managed by VMware that is set by the Storage team.
- New block data store, where we can find the VNX Unified storage available in the list.
- Constant tracking for all of the steps and tasks taking place to create any storage in vSphere through the Tasks panel.
- Build either block or NFS (file) based storage.

Features

Continuously updated configurations between Unisphere and vSphere:

- If a LUN is resident on a RAID group, the auto-Tiering features will not be available in the vSphere of this LUN (as per the RAID design) if the VNX online data migration is used to migrate this LUN to another pool rather than the RAID group from Unisphere. Once the online migration is completed, all of the properties and configurations will be updated in the vSphere, allowing storage administrators to configure any available auto-Tiering policy.
- If the LUN size presented to the VMware hosts has been fully used and users need more capacity on the pool, storage administrators can easily use the Unisphere to expand the LUN or use the online LUN migration to a bigger LUN. Once the space is added to the LUN, the storage administrator can use vSphere to increase the added space to LUNs. In minutes, the extra space requested by the users is available for them.

VNX LUN migration feature

LUN migration is managed through Unisphere. The Migrate option invokes a dialog that allows the user to select the destination and rate of the migration for the session. The migration rate and the number of concurrent migration sessions can be set to minimize performance impact. A set of CLI commands are also available for managing migrations.

One of the best benefits is online LUN migration independently or outside of the vCenter. Suppose the storage administrator created a pool that doesn't have sufficient resources that resulted in performance issues with the users. Instead of using the Storage vMotion to move all of the Virtual Machines around, we can use the VNX virtual LUN migration option between storage devices while the LUNs are available and online.

The IOmeter is software used to measure the performance impact of VNX disks/LUNs through emulating the workload on the LUNs and track the changes and effects after different types of I/O work load.

Storage administrators can set up the IOmeter software on the Virtual Machine to track the I/O traffic and performance of the Virtual Machines.

On each virtual machine in vSphere, these features are present in Unified Storage,

- VNX storage efficiency features data compression. This can be used through the Unified Storage features and can be done on a specific VM or on the entire data store.

VNX has introduced an easy solution for storage administrators that struggle with what the VMware administration is talking about. Unisphere, the storage administration console of the VNX, can be configured with the VMware console information. Once the information of the VMware vCenter is added to the Unisphere, it can view the Virtual Machines.

From Unisphere, administrators can find the relations between the LUNs, disks ,and virtual machine, and map the relation back and forth through the "Virtualization" tab in Unisphere.

VMware vSphere Storage APIs – Array Integration (VAAI) is also referred to as hardware acceleration or hardware offloads APIs. They enable the communication between ESXi hosts and storage array. The APIs define a set of "storage primitives" that enable the ESXi host to offload certain storage operations to the array. This reduces resource overhead on the ESXi hosts and can significantly improve performance for storage-intensive operations such as storage cloning and zeroing.

A storage integration feature that increases virtual machine scalability, VAAI consists of a set of APIs that enables vSphere to offload specific host operations to EMC VNX storage arrays. These are supported with VMFS and RDM volumes. VAAI 2 Includes Block enhancements and File support.

VMware Aware Storage APIs (VASA) is a set of vCenter providers that enable a VMware administrator to retrieve system capabilities for VASA-enabled storage systems. This is achieved via defined names and details associated with a "Storage Capability."

VASA consists of two dependent network services:

1. A vCenter service (VMware vSphere profile-driven Storage).
2. EMC Storage Management Initiative Standard (SMI-S) Provider service for VNX.

EMC VASA Provider Capability Description

Capabilities		
vCenter Storage Capability	**Description**	**VNX LUN Type**
Unprotected	No RAID protection	
Capacity	RAID-protected, uses capacity oriented drives such as SAS, NL-SAS, or SATA	Flare/Traditional SAS LUN
Performance	RAID-protected, uses Fibre Channel or high-end SAS drives	
Extreme Performance	RAID-Protected, uses Flash Drives	FLARE Flash Drive
Multi-Tier	RAID-Protected, uses multiple tiers of drives	Pool LUN, FAST, FAST Cache
NFS	Unsupported	File system

The EMC VASA provider must be installed and configured on a Windows system which can be the same host running vCenter, or a stand-alone system. SMI-S can be configured to manage VNX storage systems using in-band SCSI LUN or out-of-band using the CIM interface for VNX.

The Storage Viewer feature extends the vSphere Client to facilitate discovery and identification of VMAX®, VPLEX®, and VNX storage devices that are allocated to VMware ESXi hosts and virtual machines. Storage Viewer (SV) presents the underlying storage details to the virtual data center administrator, merging the data of several different storage mapping tools into a few seamless vSphere Client views. SV enables you to resolve the underlying storage of Virtual Machine File System (VMFS) and Network File System (NFS) data stores and virtual disks, as well as raw device mappings (RDM).

EMC Unified Storage Plug-in for VMware vSphere is a VMware vCenter integration feature designed to simplify storage administration of the VNX and VNXe unified storage platforms. The feature enables VMware administrators to provision new NFS and VMFS datastores, and RDM volumes directly from the vSphere Client.

VMware/EMC integration capabilities in the VNX platform are provided at no additional cost. It covers:

- VCenter-integrated provisioning (block and NAS) via the VSI 4.1 vCenter plugin + USM module.
- vCenter-integrated visibility (block and NAS) via the VSI 4.1 vCenter plugin
- Automated and simplified path management (for both NMP and PP/vE) via VSI + Path-ing module
- Simplified data store expansion
- Simplified VM-level replication using hardware-acceleration in the VNX
- Simplified VM-level compression in the VNX
- Automated ESX host registration in Unisphere (through embedded APIs in the vmkernel)
- Automated VM-storage mapping in Unisphere (through vCenter API integration)

Advanced Storage Analytics for VNX storage arrays using VMware's vCenter Operations product, Tight integration between storage hardware and the vCenter Operations monitoring suite enables best-of-breed analytics and monitoring software delivered by VMware Operations that provides in-depth storage statistics and operations monitoring necessary for optimizing storage performance and validating that storage-level SLAs are met. The new product – VNX Storage Analytics Suite and VNX Connector for VMware vCenter Operations – will be generally available later this year.

VNX integration with Recover Point

What is RecoverPoint?

RecoverPoint has introduced a new level for block protection. It handles outages or data loss situations and eliminates effects that might be experienced in the production site, enabling organizations to easily achieve RPO and RTO goals.

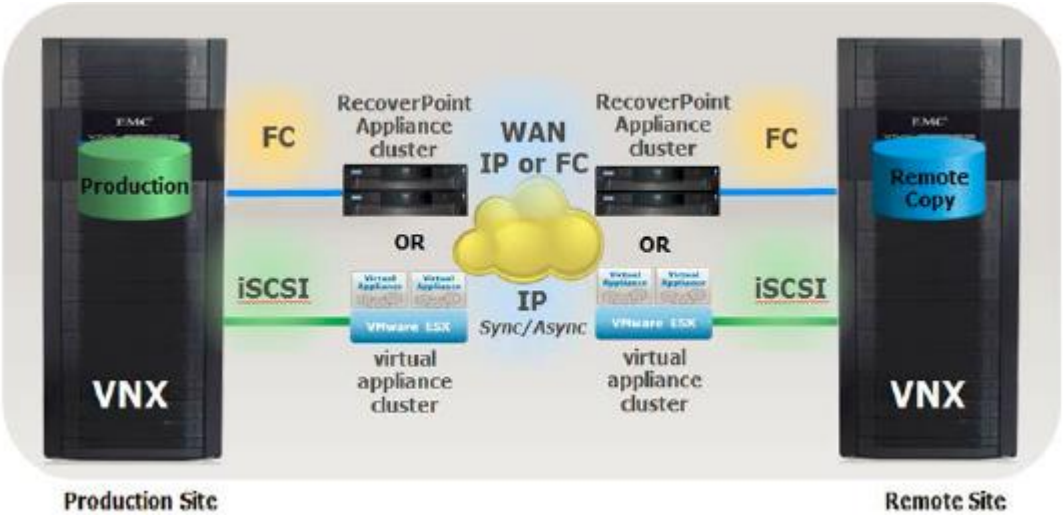
RecoverPoint also expanded its protection capabilities to include protection of Virtual Machines in VMware virtualized environments.

EMC RecoverPoint provides protection to storage array LUNs through concurrent local and remote data replication over any distance, synchronous or asynchronous, offering continuous data protection for any point in time recovery.

It supports VMAX 10K, 20K, 40K, VNX series, VPLEX, and any 3rd party array via VPLEX. Integrated with VMware Site Recovery Manager (SRM), it extends the protection capabilities of SRM beyond snapshot.

RecoverPoint Features

RecoverPoint leverages NX array snapshot capability to enhance asynchronous replication with a user-defined interval for replication. This Snap and Replicate feature adds intelligence to the asynchronous replication policy to ensure effective and efficient capture of data protection under a high data load.



Snap-based replication and VNX Operating Environment (OE) for Block 32 or later:
In RecoverPoint, snap-based replication provides point in time snaps for the production volumes. Benefits for enabling snap-based replication include reduced traffic transmission, lower RTO and, as a result, better journal utilization.

Snap-based replication requires all RecoverPoint Appliance (RPA) software involved to be running code level 4.1 or higher. The storage array must be registered on the RPA cluster through the Storage wizard. Username and password are required to register the IP addresses of both storage processors.

At the production site, VNX storage arrays are supported and they must be running VNX Operating Environment (OE) for Block 32 or latest. In addition to the RecoverPoint Splitter enabler, the VNX Snapshot enabler must be installed and active.

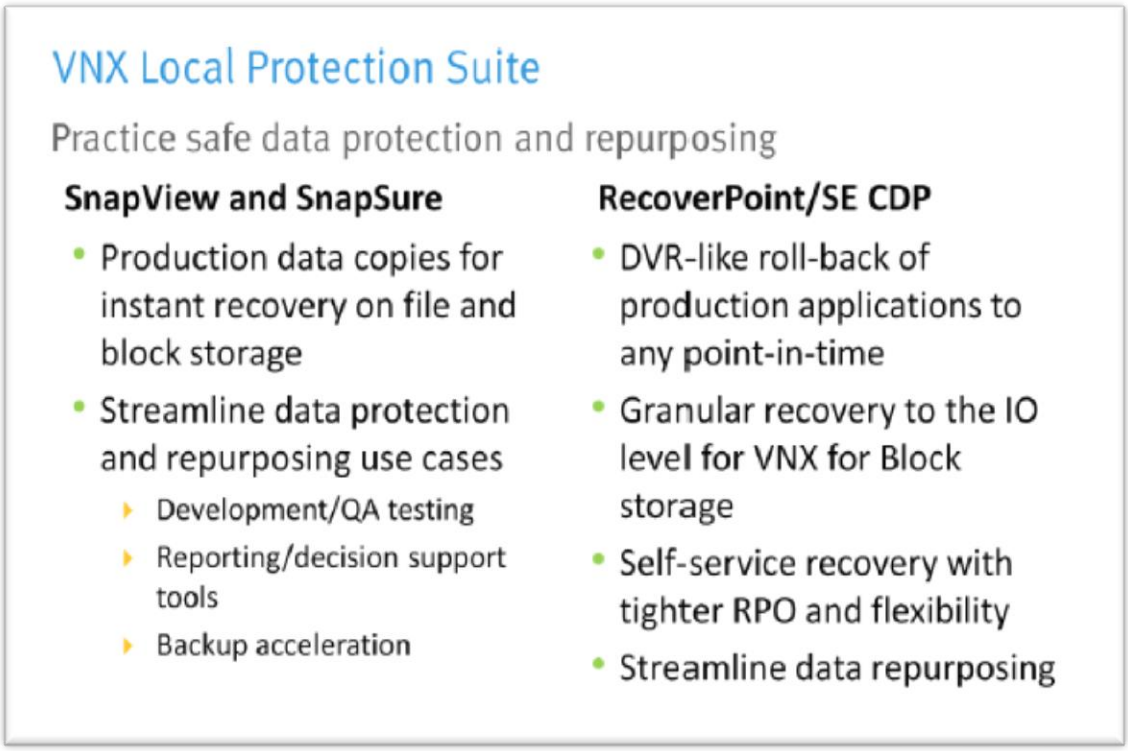
RecoverPoint also supports replication with VNX File and VNX Gateway systems for File replication. RecoverPoint replicates at a block level and provides protection for the entire NAS system. Use VNX Replicator for file system- or Virtual Data Mover-level replication granularity.

RecoverPoint Integration with Unisphere
Designed to accept plug-ins that will extend its management capabilities, Unisphere provides a plug-in for RecoverPoint from a central location.

RecoverPoint Virtual edition for VNX Series consists of RecoverPoint Appliance (vRPA) software deployed as a virtual appliance in an existing VMware ESXi VM environment. This software is currently available for the VNX series equipped with iSCSI support.

VNX local protection Suite

Local replication can significantly enhance business and technical operations by providing with access points to production data; enabling parallel-processing activities such as backups, as well as disk-based recovery after logical corruptions; and creating test environments for faster time to revenue for applications. Every business strives to increase productivity and use of its most important resource—information. This asset is the key to finding the right customers, building the right products, and offering the best service. The greater the extent to which corporate information can be shared, re-used, and exploited, the greater competitive advantage a company can gain. EMC offers local snapshots for Block (VNX SnapView) and File (VNX SnapSure) as well as a continuous Data Protection software option (RecoverPoint/SE), to provide the broadest set of native local replication capabilities in the market.



RecoverPoint/SE CRR is a comprehensive data protection solution that provides bi-directional synchronous and asynchronous replication. RecoverPoint/SE allows users to recover applications remotely to any significant point in time without impact to the production application or to ongoing replication.

VNX Total Protection Pack Is a combined data protection and recover point solution, including RecoverPoint/SE, Replication Manager, Data Protection Advisor, SnapView and SnapSure, all managed within Unisphere.

EMC RecoverPoint/SE replicating between LUNs that reside inside the same storage system or between storage VNX storage arrays provides simple replication management, lowers RPO, provides application consistent recovery, and simplifies everyday operation.

Unisphere is used to show all of the resources being used in the RecoverPoint environment.

VNX Replicator is a synchronous file system-level replication technology that complies with the customer-specified RPO. Replicator is included in the Remote Protection Suite and the Total Protection Pack for VNX systems.

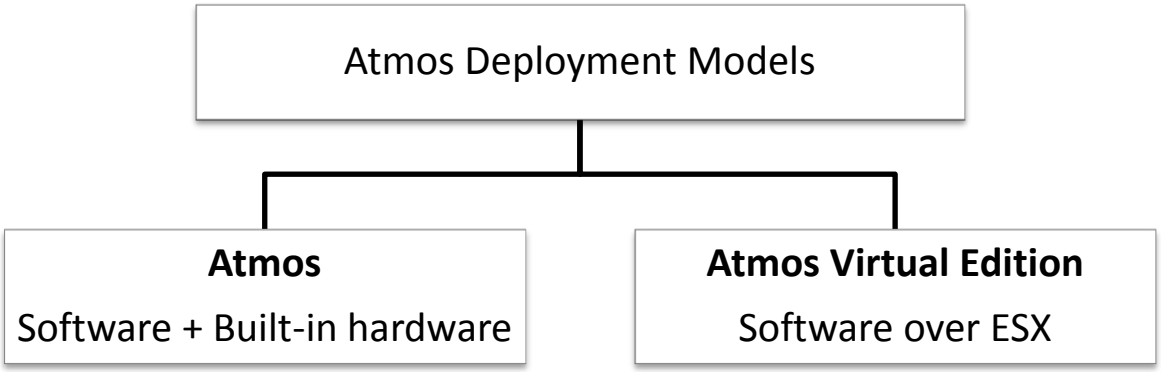
If a disastrous event occurs, VNX Replicator provides organizations with the ability to transfer the NFS/CIFS responsibilities to the DR site, eliminating impact that may have been caused to the production environment.

The DM interconnect is the communicating channel to transfer data between the source and destination. VNX replicator works by sending periodic updates to the target File system. Configure the source and target VNX systems for communication by creating a relationship between the systems and then creating the data interconnects between participating DMs.

VNX Replicator provides a manual failover capability to remedy a disaster that could make a file system or an entire VNX system unusable or unavailable. After failover, target file systems are changed from RO to RW mode. Target file systems can then be used to provide access to data. Failover may cause data loss, but it is managed by the user configuration RPO policy which is known as ‘max time of sync’ and is set to 10 minutes by default.

VNX integration with ATMOS

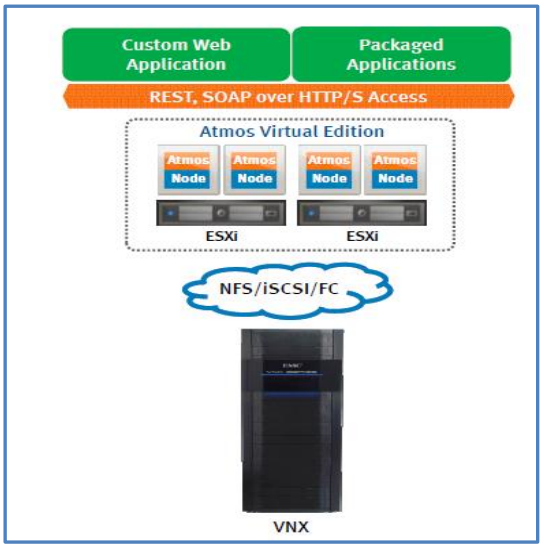
Atmos® is an object-based cloud storage platform to store, archive, and access unstructured content at scale. Using object storage architecture to manage data across the world, Atmos provides the essential building blocks for enterprises and service providers to transform to private, hybrid, and public cloud storage. .



Atmos Virtual Edition (VE) is recommended with EMC Unified Storage (VNX), expanding the storage options to Web services.

Integration with VNX

Atmos VE with VNX provides an simpler, more powerful way to build an object data store. It is easy to deploy and lowers costs as it uses already available resources on the site which will allow leveraging existing VNX features. For example, VMware supports data stores provisioned by VNX storage array to virtual machines.



Four main components to build Atoms VE over VNX Storage:

1. Storage
2. Network
3. VMware ESX servers
4. Atmos Virtual Edition

The following are best practices of each component to optimize performance.

1. **Storage Setup Best Practices**
    - EMC best practice highly recommends balancing the workload on the physical drives.
    - Create separate pool for object and non-object data.
    - Create separate pool for Atmos boot store and data/meta data store.
    - Create the storage pool using RAID group 5 as it is giving best performance with best protection.
    - Provision the storage to the network considering multipathing to avoid single point of failure.
    - Do not use VNX replication solution for replicating Object Stores.
2. **Network Setup Best Practices**
    - Separate the VM network and the storage network to avoid traffic interruption between both.
    - Use Fibre Channel connection or 10 GbE for iSCSI.
    - Consider multipathing and make sure there is a connection to both storage processors.
3. **VMware ESX servers Best Practices**
    - At least two at each site to avoid single point of failure.
    - Create boot data store for each ESX server.
    - Create data/metadata store based on how big the object store is going to be.
    - Create the Atmos nodes and evenly distribute them across the ESX servers.
    - It is recommended to use the EMC Virtual Storage Integrator (VSI) for VMware vSphere to provision and manage Atmos VE storage in VNX.
    - Maximum Atmos nodes per site is 32.
    - Maximum storage capacity per Atmos node is 30 TB.

Atmos node configuration best practice is shown below:

CPU/memory	2 vCPU, 8 - 12 GB RAM
Virtual NICs	1 private (node 2 node communication) 1 public (for end user access)
Virtual Disks	Virtual disks: - Boot disk - Meta data - Data

4. Atmos VE Setup Best Practices

- It is not recommended to use GeoParity replica for protection.
- Use Object Replication, one local using sync replica and one remote using a sync replica.
- Use Optimal Balanced for data placement options.

VNX Best Practices for Optimal Performance

VNX1

Let's start with best practices for Storage Processor Cache where it means best practices read and write Cache.

SP Cache size

It is recommended to make the following settings when configuring SP Cache on a new VNX.

- Software Enablers should be installed before setting Write and Read Cache
- Read Cache should 10% of the Available Cache where the minimum recommendation is 200 MB and Maximum is 1024
- Purpose of Read Cache is to make pre-fetching easy so it doesn't have to be large. It can be increased above the recommended values only if multiple sequential read-intensive applications will be used.
- Read Cache should be the same on both Storage Processor A and Storage Processor B
- After assigning the Read Cache, give all remaining memory to the Write Cache

SP Cache Page Size

- Cache page size determines the minimum amount of SP Memory used to serve a single I/O operation
- Default Value is 8 KB for the majority of workloads. It can be increased to the maximum value 16 KB if large block I/O size is controlling the environment

SP Cache Watermarks

- Watermarks are intended to be used as an indicator and control the flushing of the write cache
- Default Low watermark is 60% and High watermark is 80% of the Write Cache. These values can be changed in some situations such as frequent forced flushing.
- it is highly recommended to maintain about a 20% difference between them
- In cases where VNX for file will be involved, recommended values for cache settings do no change

Using Flash Drives for Fast Cache and Fast VP

Flash Drives for Fast Cache

- When Flash Drives will be used as a Fast Cache, please use all Flash drives in bus 0 enclosure 0 (0_0), up to 8 drives.
- In cases where more than 8 drives are used
  - Spread them across all available buses
  - Mirror drives within an enclosure to avoid mirroring across 0_0
- Flash Cache is disabled by default on a LUN already created on Flash Drives as it doesn't make sense

Flash Drives for Fast VP

- Spread Flash drives across all available buses
- Avoid using enclosure 0_0

Hot-sparing

Hot-spare hard drives exist to automatically replace any failed drive

- Allocate one hard drive for every 30 drives
- Type should be the same, i.e. Flash must be spare for Flash, and so on
- Hot-spare capacity should be equal to or larger than the source drive

Hard Drives

Drives are able to sustain IOPS workload based on the drive type as follows:

	NL-SAS	SAS 10K	SAS 15K	Flash
IOPS	90	140	180	3500

Bandwidth assumes large-block sequential with multiple streams.

- Spinning drives provide about 30 MB/s of bandwidth capability; Flash drives provide about 100 MB/s
- For bandwidth with spinning drives, VNX models scale up to the sweet spot drive count, which varies by model as per the following table

	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Drive count	60	80	120	140	300

Fast VP

- It is recommended to maintain some unallocated capacity within the pool to help with relocation schedules when you use FAST VP
- Relocation will reclaim 10% free per tier. This space will be used to optimize relocation operations but also assists when new LUNs are being created that want to use the higher tiers
- Enable Fast VP on a pool even if it has only one tier to provide ongoing load balancing of LUNs across available drives
- By default, a VNX for a file system-defined storage pool is created for every VNX for Block storage pool that contains LUNs available to file, called a Mapped Storage Pool. All LUNs in a given file storage pool should have the same FAST VP Tiering Policy. If different tiering policies are preferred or required, you can create a user-defined storage pool to isolate File LUNs from the same Block storage pool.

Fast Cache

Fast Cache is a very important feature in VNX that improves performance. Fast Cache enabler must be installed to enable Fast Cache. Misconfiguration of Feast Cache might impact performance. It is strongly recommended to consult EMC Professional Services to provide the best design for your array.

- For VNX1

  Fast Cache drives must have the same capacity.

  Always use DAE 0 on every bus to avoid latency.

  For the highest level of availability, mirror Fast Cache drives between buses. For example, the primary disk is 1_0_0 and its secondary is 2_0_0.

  Upgrade to the latest patch of code 32 is highly recommended.

- General recommendations:
  - It is not recommended to enable Fast Cache on LUNs with small sequential I/Os like:
    i. Reserved LUN Pool RLP
    ii. Clone private LUNs CPL
    iii. RecoverPoint Journal LUNs
    iv. Write Intent Logs
    v. SnapView clones
    vi. MirrorView Secondary Mirrors

- Fast cache drives ratio is 25:1. For every 25 drives that are in LUNs or pools with Fast Cache enabled, at least one drive should be added to Fast Cache.
    - Minimum recommended drive count for each VNX model is shown below.

Model	Number of drives
VNX5200	2 drives
VNX5400	4 drives
VNX5600	4 drives
VNX5800	8 drives
VNX7600	8 drives
VNX5100	2 drives
VNX5300	4 drives
VNX5500	4 drives
VNX5700 and VNX7500	8 drives

- From the CPU Utilization perspective
    - Less than 60% SP CPU Utilization; enable Fast Cache groups of LUNs or one pool at a time. Let them equalize in the cache and ensure that SP Utilization is still OK before turning on Fast Cache for extra LUNs or pools.
    - From 60-80% SP CPU Utilization; scale in carefully; enable Fast Cache on one or two LUNs at a time, and verify that SP CPU Utilization doesn't exceed 80%
    - More than 80%; don't activate Fast Cache
- Avoid enabling Fast Cache for a group of LUNs where the collective LUN capacity will exceed 20 times the total Fast Cache capacity
- Since Fast Cache is a pool-wide feature, enable/disable at the pool level for all LUNs in the pool

- Avoid enabling Fast Cache for LUNs including SavVol because SavVol Activity tends to be small block and sequential, features that do not benefit from Fast Cache. When enabling Fast Cache on a Block Storage Pool with File LUNs, use a separate storage pool or traditional RAID group for SavVol Storage.

iSCSI consideration for the Management IPs

It is not recommended to use IP addresses of the same subnet for the management LAN of the SP and the iSCSI HBAs. Having them in the same subnet may cause a short interruption.

VMware

When VNX integrates with VMware; a LUN in multiple Storage Groups should have the same HLU number to avoid reservation conflict.

VNX2

Storage Processor Cache

A new feature for this VNX generation, VNX2 automatically assigns the storage processor memory. There is no need to configure Write Cache, watermarks or any other parameter related to the storage processor cache.

Hot-Sparing

- Beginning with VNX Block Level 05.33, any unbound drive can be considered for sparing, providing high flexibility
- It is recommended to plan one additional drive for every 30 drives
- Distribute unbound drives across available buses
- Make sure that unbound drives for each drive type are available and capacity is equal to or larger than the provisioned drives

Hard Drives

- Match the suitable drive type with the expected workload as per the following
  - Flash drives for extreme performance
  - SAS for general performance tier
  - NL-SAS for less active data where it is recommended for streaming data, aging data, archive, and backups

Traditional RAID Group Configuration

- From the Drive count perspective
    - For RAID Group 5, use 4+1 or 8+1
    - For RAID Group 6, use 8+2
    - For RAID Group 1/0, use 4+4

Storage Pool Configuration

- Use think LUNs for highest performance
- Don’t trespass Pool LUN - use migration if required
- Use separate pool for different storage profiles required
- DON’T enable FAST Cache for small block sequential or large block I/Os.

File Storage Configuration

- Create dedicated pool for File
- Only use thick LUNs
- If virtual provisioning is required, use thin-enabled file system.
- Create all LUNs with the same size
- Make pool utilization less than 95%

Fast Cache

- Distribute Fast Cache drives evenly across the available backend buses
- Maximum of 8 Fast Cache drives per bus includes the hot spare
- Always use DAE 0 on every bus to avoid latency .

Author Biographies

Omar Aboulfotoh

Technical Support Engineer

Omar is a Technical Support Engineer in Unified Storage Division, Global Technical Support, EMC Egypt Center of Excellence (CoE). He has been in the IT industry for 3 years.

He holds a degree in Electronics and Communications Engineering and is studying a Master's program at Faculty of Engineering, Cairo University.

He is a certified EMC Proven Professional Specialist in VNX and SAN. Also, he is VCP- and vCloud-certified.

Wessam Seleem

Technical Support Engineer

Wessam is a member of the Unified Storage Division Recovery team. Wessam has more than 10 years' experience in the IT field. Wessam joined EMC in 2011.

Wessam holds numerous certifications in VNX, Microsoft, and RedHat.

Maha Ate**f**

Technical Support Engineer

Maha is a Unified Technical Support Engineer in the Unified Storage Division. She holds a degree in Computer and Systems Engineering from Ain Shams University.

Maha holds several certifications - VNX Specialist and Implementation, SAN Implementation, VMware Administration, RecoverPoint Specialist, and VCE Associate.

Mohammed Hashem

Technical Support Engineer

Mohammed is an EMC Technical Support Engineer, Unified Storage Division. He is an EMC Proven Professional Specialist in VNX, SAN Management, and RecoverPoint.

Mohammed holds a Bachelor's degree in Communication Engineer, German University in Cairo.

Appendix

- http://www.clipper.com/research/TCG2011008.pdf
- EMC Knowledge Base Articles
- http://vblog.matt-taylor.org/2010/11/07/fast-and-fast-cache-some-quick-highlights/
- https://support.emc.com/docu45636_Atmos-Planning-and-Considerations-Guide.pdf?language=en_US
- https://support.emc.com/docu44317_Atmos-Release-Notes.pdf?language=en_US
- https://support.emc.com/docu53618_Atmos-2.2.0-Administrator's-Guide.pdf?language=en_US
- https://www.emc.com/collateral/software/white-papers/h9505-emc-atmos-archit-wp.pdf
- http://www.emc.com/collateral/hardware/white-papers/h8217-introduction-vnx-wp.pdf
- https://www.emc.com/auth/collateral/hardware/data-sheets/h8520-vnx-family-ds.pdf
- https://support.emc.com/docu42149_White-Paper:-VNX-Virtual-Provisioning-VNX5100,-5300,-5500,-5700,-and-7500-Applied-Technology-.pdf?language=en_US
- https://support.emc.com/docu53296_VNX5100,-VNX5150,-VNX5300,-and-VNX5500-Global-Services-Product-Support-Bulletin.pdf?language=en_US
- https://support.emc.com/docu42151_White-Paper:-Unisphere:Unified-Storage-Management-Solution-VNX5100,-5300,-5500,-5700,-and-7500-.pdf?language=en_US
- http://virtualgeek.typepad.com/virtual_geek/2011/04/vsphere-vnx-integrated-howto.html
- https://www.emc.com/collateral/hardware/technical-documentation/h8229-vnx-vmware-tb.pdf
- https://www.emc.com/collateral/software/15-min-guide/h8541-vnx-vmware-top-3.pdf
- https://www.emc.com/collateral/demos/microsites/mediaplayer-video/emc-recoverpoint-se-emc-unisphere-emc-vnx-series.htm
- http://www.emc.com/collateral/software/data-sheet/h2769-recoverpoint-ds.pdf
- https://www.emc.com/collateral/white-papers/h12079-vnx-replication-technologies-overview-wp.pdf

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO RESPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.